University of California
Lawrence Livermore
National Laboratory

Livermore Computing

Delivering TeraScale Computing

Mission

The Center, Livermore Computing's Production Computing and Advanced Development Division, provides leading-edge computational infrastructure to support the development and application of science-of-scale calculations targeted at pressing national issues. In collaboration with other Department of Energy (DOE) laboratories and universities, the Center engages in research to expedite the development of a usable terascale computing environment. The Center provides access to stockpile stewardship research, which alone requires an enormous computing system. However, because LLNL derives much of its strength from its multiprogrammatic and multidisciplinary base, the Center is dedicated to serving the aggregate computation needs of its research community, whose requirements extend from the midscale to the terascale.

Impact

The Center is dedicated to expediting the research scientist's productivity. The sheer magnitude of the scientific problems that LLNL scientists must now address will drive the Center's research and development work. The Center is committed to the goal of a globally scalable network-based computing environment that, in its maturity, will bring the huge capability of the Center onto the scientist's desktop.

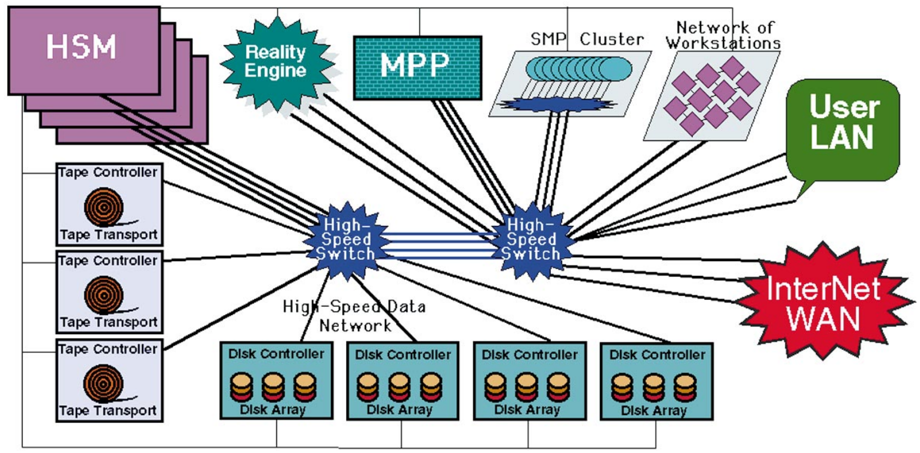

Globally scalable computing environment.

Dramatic world changes are driving the evolution of LLNL's mission, which encompasses national security, energy resources, and environmental quality. The Center's primary mission is to provide a computing environment for the Department of Energy's Stockpile Stewardship programs, in particular for the Accelerated Strategic Computing Initiative (ASCI) program, which recently placed a very large IBM parallel system at LLNL (see The ASCI Blue Pacific Teraflop Platform). The Center is also chartered to serve all programs and scientists engaged in research at LLNL, which is interested in leveraging the ASCI investments to serve its extensive multidisciplinary base.

Hardware Architecture

Movement toward a distributed computing environment utilizing multiple teraflop computer platforms, tens of terabytes of disk storage, petabytes of archival storage, and copious ancillary equipment (e.g., remote rendering servers and smaller symmetric multiprocessor-based [SMP] computer platforms) dictates a network architecture very different from what has been traditional. In the past, each device, such as the mainframe computer, represented a relatively complete resource—an island unto itself, with its own local disk and its own file system. All resources shared a bus-based network.

This inflexible picture has been replaced with the prospect of a scalable network, or the "network is the computer" model. In this new picture, every device is connected to a very high speed switch-based network.

Software Architecture

The software environment to which the Center is moving is in partnership with Sandia National Laboratory and Los Alamos National Laboratory through the ASCI program's Problem Solving Environment (PSE) research and development. The goal of this Tri-Lab collaboration is to assure that the "network is the computer" model works effectively and presents a unified view to the research scientist.

The plan is for a layered software environment. It leverages applications from commercial, public domain, and Tri-Lab developers to support secure, distributed high-performance computing in the areas of parallel I/O, mass storage systems, scientific data and information management, visualization, and resource allocation and management.

In such an environment, for instance, it is essential that the secondary and tertiary storage devices

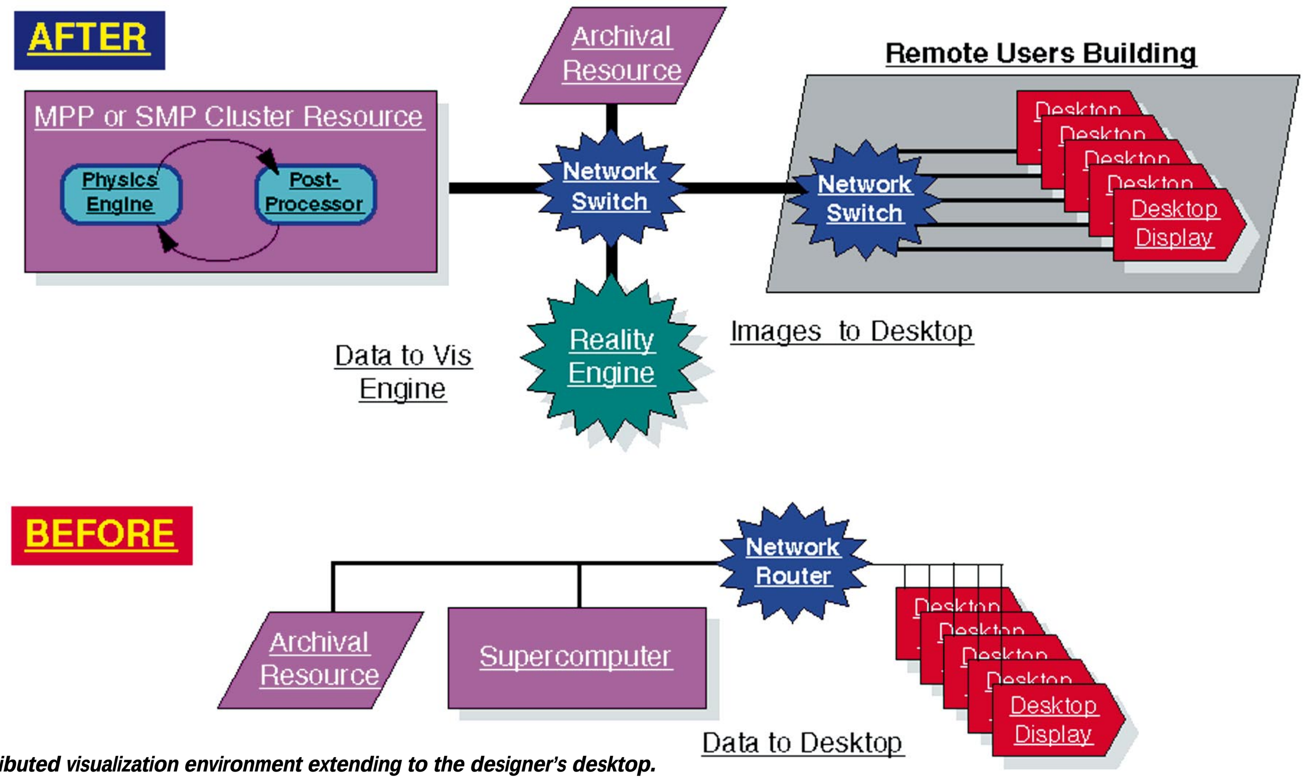

Distributed visualization environment extending to the designer's desktop.

be accessible in parallel at very high aggregate bandwidths.

In another example of the network model, consider the nature of the visualization environment: A scientist could read stored terabyte data sets at high speed into the terascale computer for postprocessing, and move a smaller postprocessed data set to a powerful rendering platform, which then sends the images to the remote desktop. Such are the kinds of solutions that must accompany any superteraflop computer to assure efficient use of the resource.

The ASCI Blue Pacific Teraflop Platform

LLNL has recently taken delivery of a 512-node SP2 parallel processor, as a first step in the eventual delivery of a very large 3-Tflop SMP computer platform, called ASCI Blue Pacific. A sequence of refresh steps in 1997, 1998, and 1999 will move the system from single processor nodes, to 4-way, then to 8-way SMP nodes.

By 1999, the system will consist of a 512-node, 8-way SMP configuration with 2.5 TB of memory, 75 TB of disk, and a peak speed of 3.2 Tflops. Most nodes will realize 3 GB of memory, but 64 nodes will feature 16 GB of memory. This creates a very flexible architecture for the execution of mid-scale to extremely large 3D applications. This resource will be able to securely shift from classified to unclassified use in a matter of hours, allowing scientists the benefit of collaborative research with workers in industry and academia.

A Multiprogrammatic, Institutional View of Computing

The needs of the stockpile stewardship mission at LLNL have driven the development of terascale computing capability. However, because LLNL derives its strength from its strong multidisciplinary base of research scientists, the Laboratory has allocated resources to allow multiple programs and disciplines to benefit from the very high end computational capability at the Center. Livermore Computing is currently building a powerful cluster of symmetric multiprocessors for this purpose.

As the IBM system evolves into a resource of surpassing capability, LLNL may elect to associate with the ASCI Blue Pacific computer a similar system that sits perpetually in the unclassified environment. The unclassified system would be used for code development and smaller production runs. Scientists would move their applications onto the ASCI Blue Pacific platform during the unclassified part of its duty cycle. Such a strategy would allow a variety of disciplines at LLNL to benefit from the large-scale modeling made possible by the ASCI Blue Pacific computer.

Our Future

The Center's charter is to provide innovative, high-impact technology that enhances the nation's scientific and technological strengths. Developing, integrating, and applying this technology in support of a broad spectrum of national goals requires new forms of cooperation among the national laboratories, universities, and industry.

The ASCI partnership is a part of this growing collaborative picture, which could also include National Science Foundation (NSF) centers and universities.

For more information, contact Michel McCoy, 510-422-4021, mgmccoy @llnl.gov.

Work performed under the auspices of the U.S. DOE by LLNL under contract No. W-7405-ENG-48